it has a guaranteed learnability. There are indications, also, that the number of ways of learning (hence, from Section 14, the resilience that might be acquired by a student learning in this conversational domain) depends upon the quality (Section 11) of the underlying entailment mesh.

On the debit side of the account, this entailment structure must not be violated by seemingly innocent manoeuvres, such as discarding arbitrary parts of it. As a rule, such infringements are liable to loose cyclicity and/or consistency though judged by looser criteria of course material organisation (those criticised in Section 12) the entailment structure is often quite flexible.

According to the present argument, this is a comment on the nature of knowables. True, almost anything may be learned in some way. You can learn;

Bin, Zin, Bish, Pish

or even

* Z # !

or any other nonsensical combination of tokens. But you do so by imposing your own pattern upon them (perhaps out of kilter with the meaning, if any, I had in mind when writing the tokens down). Someone who claims to teach (or some heuristic that is said to do so) must lock onto pattern in a tutorial conversation. We insist that the pattern in question belongs to the entailment structure viewed as a grammar like entity which bounds the class of grammatical (or legal) tutorial transactions, and insist, as well, that any legal pattern is consistent and cyclic. As a final (but probably crucial) refinement, structure should represent as many legal (hence, consistent and cyclic) patterns as possible.

Chapter 8. A Currently Mechanised Form of the Course Assembly Heuristic

The interchange between the source and the AI, described in the last chapter, is governed by a heuristic and is associated with special algorithms for testing cyclicity, checking consistency, pruning a given entailment mesh and so on. As described, both the source and the AI are human beings, of whom the AI is responsible for arbitration (if needs be for enforcing the heuristic) for demanding the (in practice, mechanised) routines that execute the special algorithms, and for displaying the results of these computations to the source. This chapter contains a brief account of a mechanised version of the AI which is currently operating as an evolutionary B heuristic (the routine is listed, with full annotation, in Appendix K). It operates in the context of the CASTE system (under the CET heuristic listed and described in Chapters 4 and 5) and it is an extension of the CET heuristic. But it is not, of course, confined to a particular subject matter, such as probability theory, and though it can be used with a modelling facility as restricted as STATLAB, it is generally attached to a much more liberal facility. For example, any of the modelling facilities to be discussed in the next chapter is acceptable, and the programming language, TELCOMP is actually employed as the present modelling facility (i.e. a "model" is any TELCOMP program capable of execution within certain restrictions upon "statement" number and storage).

1. Outline

The source (A) starts out as a student who has learned the topic relations in an existing conversational domain under the strict A, B conversation maintained by the CET heuristic. Let this domain be R; and so it remains whilst A occupies the role of student under the course assembly heuristic. But A is allowed, in certain conditions, to opt out of the student role and to act, instead, as a subject matter expert who may extend the domain. If the change of role takes place at occasion n, then R becomes R(n) (as promised in Chapter 5). Further, at the end of occasion, n, fresh topic relations may have been added to the domain and if so R(n+1) is larger than and different from R(n); that is, the domain evolves. The word "user" tags either of A's roles, student or source.

2. Preliminaries

Certain features of the system, which have already been mentioned from time to time, needs to be mustered together in order to explain this process.

2.1. As $D^1(R)$ is defined, it contains an operator data base associated with each node and thus exhibits the cyclic connections suppressed by pruning i.e. with each $D^1(R)$ there is a cyclic mesh Ω. This is stored in ES(R) (see the listing in the Appendix) but only the hierarchy-like L^1 description is displayed if the user is a student. However, the entire structure can be displayed, on request, to a source.

2.2. The routine TagAim under an aim and a goal statement G (the listing of the CET heuristic in Chapter 4 Appendix D and Appendix I) either leads to the specification of a workset in which for at least one kernel k all nodes in Im Ent Set k, G are marked already as being understood or, if the student opts to explain a topic relation, R_i that is not currently marked as understood then this explanation is accepted (and R_i is subsequently marked as understood) if and only if he gives a correct explanation of R_i and if for at least one kernel k (which he is free to cite) all nodes in at least one Im Ent Set k,i can also be correctly explained (in each case, explanation involves building a model in the modelling facility proper to R_i).

2.3. Since the fine structure family of L^1 descriptors is not full it is possible for a user to point at certain place holders (or L^1 classes) in the entailment structure that are not, currently, occupied by a node which has an associated topic relation. (If treated as a student ostension, the system responds to this transaction by a statement that the node is blank or void).

2.4. In Chapter 6 and Chapter 7 we noted the possibility that L^1 descriptors can be added to those already existing, by means of the repertory grid technique, for eliciting constructs (here regarded as "attributes" or "properties" named by unary but

many valued predicates). It is also possible to evoke constructs that are topic relations (noted earlier in the last chapter as part of the unzipping operation). The fundamental conversational forms are shown in Icon 15, in Icon 16, and in Icon 17 of Chapter 6; all of which are readily mechanised. Descriptor list extension (constructs named by unary but many valued predicates) is instrumented by the descriptor selection routine of the course assembly heuristic (in the Appendix). The citation of a new topic relation leads to a process which, in a strict conversation, images the unzip operation noted in the last chapter. In fact a new topic relation is treated as an aim node which is also a goal node and is subject to the TagAim process. Only, in this case, there is no model for the newly named topic relation and one must be constructed by the user within the modelling facility. Nor is there any Im Ent Set for the newly named node: and one must be constructed by the user (that is, the newly named topic relation must be unzipped). There are several possibilities:

(a) The members of the newly named node's Im Ent Set may be topic relations specified in the existing entailment structure. If so, these are subject to TagAim and must be explained in terms of the existing structure.

(b) The members of the newly named nodes Im Ent Set consist in fresh primitives. If so, these must be explained by modelling and the models must be compatible with models for the existing topic relations, in the sense that each fresh primitive must belong to the Im Ent Set of at least one existing node z (this relation must be specified). Further the model for the fresh primitive must form part (under the user stipulated relationship), of a model for R_z.

2.5. As in Icon 15, Icon 16 and Icon 17 of Chapter 6, the addition of a topic relation to the existing conversational domain is represented by a parametric arrow penetrating the positions in the conversational skeleton reserved for the descriptors, $D^1(R)$ and/or for $D^0(R)$. Any parametric operation is interpreted as the inscription of data in a storage locus in an appropriate department of the interface. Thus R(n) replaces R in each compartment of the skeleton, so that the descriptions become $\langle D^1(R(n)), D^0(R(n))\rangle$.

2.6 *Prerequisites for operation.* If pressed to justify the course assembly process described in the last chapter (and thus to resolve the theoretically crucial issue of how a conversational domain is specified) we should turn to the CASTE course assembly or evolutionary heuristic described in the present chapter. This would not, in general, be a practical recommendation (for some materials, it is) but a means of clarifying what really goes on. It is thus interesting to summarise the prerequisites for applying this heuristic. They are as follows:

(a) At least two L^1 descriptors must be understood (they are not nodes in R but the descriptors in $D^1(R)$).

(b) At least five $\underline{Prim}^1$ must be understood (more or less, depending upon the choice of a closed set of relational operators).

(c) For n = 0 and R(o) there is an L description $< D^1(R), D^0(R) >$ which is an entailment structure (it may be very small; for example, one analogy relation is quite sufficient).

(d) All $\underline{Prim}^0$ in R(o) must be understood; hence, a modelling facility can be specified or L^0 can be given an interpretation with respect to some (quite arbitrary) field.

3. The Heuristic

Using this notation, the transactions involved in executing the course assembly heuristic are shown in Icon 18. The underlying CET heuristic is the support S^1_B, S^0_B, appearing in Icon 9. The support (realising the "cognitive reflector" as before) is both constrained and to some degree augmented by the routine listed in Appendix K (depicted in Icon 18 as C^1, C^0). Thus, the user engages in conversation with a composite system represented in Icon 18 as C^1, S^1_B : C^0, S^0_B and this conversation has a potentially variable domain R(n).

3.1. *Overview.* The user, A, may opt, within limits, to assume the role of student or source (subject matter expert, or innovator). The basic role, which is enforced under certain conditions, is student but the role of source can be adopted if the student points to a place holder (class defined in L^1 descriptors) that is not occupied by a node and if he rejects the student oriented response "blank" or "void"; choosing, instead, to assert the existence of a new node, to give it a name, and to cite a topic relation that corresponds to this name.

In entering the role of source, the user subscribes to a different experimental contract; subserving the norms considered in the last chapter (so, for example, he must add his new topic relation to the

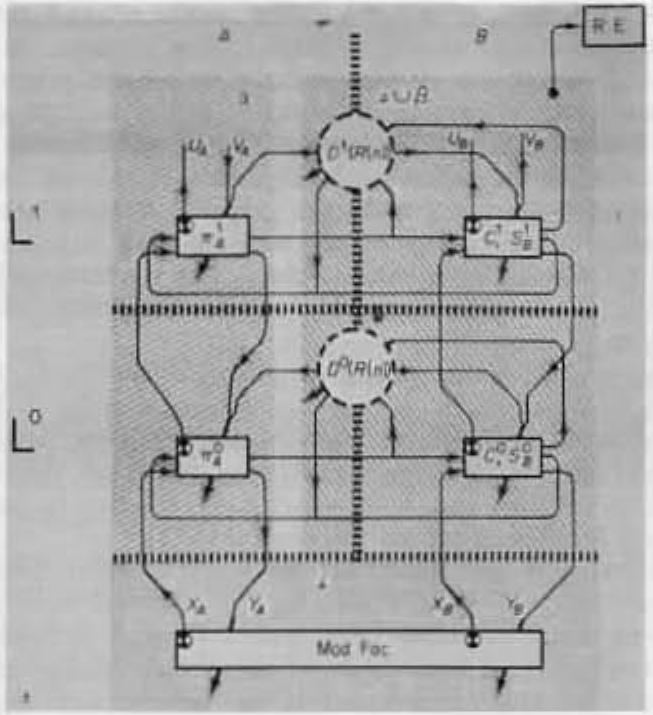

Icon 18.

existing conversational domain in such a way that the resulting conversational domain is cyclic, consistent, and fully described).

3.2. The following account conveys the gist of the operation of the course assembly heuristic.

(1) Let A have the role of "student" and at occasion n ostend (by L^1 descriptor indices) a class of place holders that currently has no member ($< E, 2, *, \phi >$ for example, in the entailment structure for probability theory).

(2) Let A reject the blank or "void" reply and thus opt to enter a fresh role "subject matter expert".

(3) B asks A what he would call a node, thus described, if, in fact, it existed.

(4) A gives a name (it is checked for uniqueness, to secure consistency, but is generally accepted by B). Let A's name be "Twinkle". An external observer's name is $i_n(max) + 1$; and it acts as a temporary or tentative entry. If A's construction is successful

the new node becomes i_{n+1}(max) at occasion n+1. Permanent entry takes place provided A can do as follows:

(a) Form or program at least one working model that brings about whatever relation he calls "Twinkle" in the modelling facility.

(b) Given a display of the entailment mesh (not just of the pruned and described version of it which is output by B) relate the node of "Twinkle" to the nodes of one or more other relations $\{R_i\}$, $\{R_j\}$ (using any relational operators for this purpose) such that $\{R_i\}$ is all superordinate to "Twinkle" and $\{R_j\}$ is all subordinate to "Twinkle".

(c) If the node of "Twinkle" is superordinate to *Head* in R(n) then a primary L^1 descriptor is elicited (as in the last chapter) and B prunes the entire entailment mesh under the new head ("Twinkle") and the new or augmented primary descriptor. If the mesh is cyclic under this head, B accepts the construction up to this point; if not, B rejects it and returns control to (b).

(d) If "Twinkle" is not superordinate to Head R(n) then cyclicity is checked by B without eliciting a primary descriptor, unless it happens that $\{R_j\}$ is/are primitives; if so, each R_j must be modelled and related as subordinate to some topic relation in R(n); after which cyclicity is checked, as before, and the construction is either accepted by B or rejected (in the latter case control is returned to (b)).

(e) Further L^1 descriptors are requested by B but only demanded if, as a result of instating the node of "Twinkle", the existing fine structure family of L^1 descriptors would become full. In either case, L^1 descriptors are elicited as constructs using an appropriate variant of the repertory grid method.

(f) For each fresh L^1 descriptor, A is required to state its values for "Twinkle" and for all R_i in R.

(g) Under the relations asserted, the node (or nodes) of R_j belong to the <u>Im Ent Set</u> i_n (max) + 1 of "Twinkle". The model for bringing about "Twinkle" (produced in (a)) is either compatible (under the asserted relations) with the existing model or models for the R_j of all nodes in sofar as at least one of the <u>Im Ent Set</u> k i_n (max) + 1 that may be asserted are all parts of the model for "Twinkle") or else it is incompatible. In the former case, the model for "Twinkle" is instated as D^0 (Twinkle) = $D^0(Ri_n$ (max) + 1), and becomes part of D^0(R(n+1)). If there is an incompatibility, control is returned to (a) and the form of incompatibility is exhibited.

(h) B tells A that he must momentarily revert to a student role and that A must learn "Twinkle", by progressing through its asserted <u>Im Ent Set</u> (which is trivial only in the improbable circumstance that all these nodes are marked as being understood) and to learn $\{R_i\}$ in the putative $R_{(n+1)}$, using all of the <u>Im Ent Set</u> k, l to which the node of "Twinkle" belongs. If A succeeds in the student role then the relation "Twinkle" is established in D^1(R(n+1)) and R(n) → R(n+1).

3.3. The process, though cumbersome, has no end. If A innovates, the mesh is pruned under A's own head node and this process eventually leads to the invention ((c) or (e)) of descriptors that are applied to all the existing nodes. These, in turn, have values that (in conjunction) fail to specify an existing node and fresh descriptors are furnished.

3.4. One defect of the existing heuristic is due to technical considerations only (those discussed in the last chapter; it is necessary to carry out some manual operations in order to instate and describe a fresh node). These difficulties are irrelevant to the system under construction, in which $D^1(R) \equiv ES(R)$ is represented on display tubes.

4. More Basic Limitations

Two fundamental restrictions are lifted, incidentally, by the technical modifications.

The present mode of evolution is unwieldy because, on account of search time (in the computer representation of D^1(R)) it is necessary to confine either tutorial or evolutionary operation to one and only one aim node. This constraint is severe. It does (as it is meant to) restrict TagAim to a fairly simple form, in which the index list is finite and ramifies over search trees. For several aims, the TagAim procedure must generate an indefinite number of indices that may form cycles (like the existing goal search routine when it deals with relations that are analogies, for example). This structure is not pathological (it exists, in a bounded form, in the existing routine and is similar to Weston's (1973) system for cognitive data base manipulation) and it opens up the intuitively plausible possibility of pointing at clusters of place holders that may either be identified with one new topic relation or several new topic relations to be dealt with in one act of innovation.

The other restriction is that A is bound to model a new topic relation (Stage (a) in Section 3.2) in one modelling facility. The restriction is again inessential. The most damaging consequence of the constraint is that A is discouraged from inventing topic relations (perhaps the most important kind) that are analogies between different parts of a heterogeneous field. Such inventions are possible, by dint of setting up special L^1 descriptors and/or L^0 primitives in an (otherwise homogeneous) modelling facility and this ritual can be performed in the system as it stands. But it is unnecessarily tedious and is circumvented in the system under construction, where it will be possible to make several independent models on the same occasion, n, and to execute them with or without interaction.

If a source's thesis about a subject matter satisfies the consistency and cyclicity requirements, then an entailment structure can be derived by the methods discussed in the last two chapters. Given this structure there are simple algorithms for generating all the multiple choice or list questions (the $\underline{PQuest^0}$s) that satisfy any given specification; in particular such questions are defined for each topic relation R_i in the conversational domain R, from which the entailment structure is derived. For each R_i it is also possible to issue a base command to bring about R_i in some or any way; it is part and parcel of the construction procedure that the source can obey such a command, usually by making a model which, on execution, brings about R_i. So far as the source is concerned, the model can be made in the abstract; as an explanation of R_i. But, unless we are prepared to interpret a student's explanations (by analysis of L^0 expressions) the source must furnish a standard modelling facility or, more often, several of them appropriate to different parts of the conversational domain. In fact, the act of describing the models built in a modelling facility is (usually, in our experience), the easiest way of interpreting explanations, though other methods of interpretation are possible. The non-verbal explanations of R_i that are counted as legal according to the source's thesis, constitute some (and without special qualification all) of the models that can be built in the facility proper to R_i and that, if executed, bring about, maintain, or satisfy R_i.

This chapter is concerned with modelling facilities. As a first step (Section 1), it is useful to consider a slightly more general category of environment, namely *laboratories;* (meaning workshops, construction shops, studios, demonstration rooms in addition to special scientific laboratories). Laboratories are familiar and an analysis of the activities that go on in them illuminates the nature of modelling as well as casting some scattered light on the modelling that might go on in facilities slightly more liberal than those that currently exist. It is useful to suppose, throughout this part of the chapter, that any operation carried out in a laboratory is sensed electrically (like the plugging